

Issues in Multiparty Dialogues

David Traum

Institute for Creative Technologies
University of Southern California
Marina del Rey, CA, USA

Abstract. This article examines some of the issues in representation of , processing, and automated agent participation in natural language dialogue, considering expansion from two-party dialogue to multi-party dialogue. These issues include Some regarding the roles agents play in dialogue, interactive factors, and content management factors.

Most formal and computational studies of natural language dialogue have considered only the two-party case. E.g., communication between two people, a person and a dialogue system, or a pair of agents. In this article, we consider several issues in dialogue management, and how the nature of the problem changes when considering multiple participants. For many of these issues, we refer to the dialogue models in the *Mission Rehearsal Exercise* (MRE) Project [1, 2]. The MRE project [3] uses virtual humans to help train decision-making in a team context, by allowing a human trainee to rehearse simulated missions, interacting with the virtual humans using spoken and multi-modal communication in an embodied virtual world. Each virtual human maintains its own model of a plan, goals, beliefs, team tasks, dialogue state, negotiation state [4], and emotional state [5]. Virtual humans can understand and talk to the human trainee, as well as other virtual humans (using an agent communication language modelled on the physical performance of speech, indicating the verbal and non-verbal information expressed and the timing of actions). In the initial, Bosnia scenario, the trainee plays the role of an Army Lieutenant platoon leader, facing a dilemma in a peacekeeping situation. The Lieutenant must communicate with a Sergeant, a Medic, and others including platoon members and local citizens as well as more distant units by radio. Since the trainee has considerable flexibility in how he chooses to communicate, and the aim is to immerse the user in a realistic simulation, many issues in multi-party and multi-modal communication must be addressed.

1 Participant Roles

There are a number of different types of participant roles that are important for dialogue interaction. These include both local roles that shift during the conversation, such as speaker and hearer, roles tied to the activities that the dialogue is a part of, and more permanent social roles that transcend particular dialogues.

1.1 Conversational Roles

At the most immediate level, there are the conversational roles. For two party dialogue, there are the basic roles of speaker and listener/addressee. When we consider multi-party communication, there are two related sub-issues: who can receive (is intended to receive) an utterance, and who is it addressed to. For instance an agent A might want to ask a question of agent B , but might also want C to hear the question as well. Likewise, D might also hear the question even though A had no intention for D to do so. There are a number of types of other listener roles, including ratified by the speaker (intended to hear the communication) or not, known to be listening by the speaker, or not. Clark gives a taxonomy of some of these listener roles [6]. An additional consideration is whether the listener is in-context or out of context. An in-context listener (who has heard the previous relevant utterances) may interpret an utterance quite differently from one who comes in without this context (or worse, with a partial or different context). There are also roles that we can use to characterize agents with respect to a whole conversation, as well as a specific utterance. *Active Participants* may take up speaker and addressee roles in a conversation, and generally are engaged and attentive to the conversation. Overhearers (who may be ratified or not) are also part of the conversation, in that they will receive and interpret the constituent utterances, and utterances may be planned with them in mind (either to facilitate or block understanding), but do not play a main part in the conversation. Finally some agents may be un-involved in the conversation.

1.2 Speaker Identification

In two-party dialogue, speaker identification is not a real issue - any speech that does not come from oneself must come from the other participant. In multi-party situations, it may not be so trivial [7]. If just a single audio stream is present, one can use a number of features as evidence for identifying speakers. These include acoustic features of the voice itself, as well as stylistic features, and self-identifications (in the case where one can trust the speaker to provide accurate information). If multi-modal information is available, additional cues can be used. E.g., stereo microphone arrays can localize the position of the speech, and thus give clues as to the speaker's identity. Likewise, visual information (e.g., of lips moving or other speech-related gestures), can help an agent identify the speaker. When multiple agents are involved in dialogue, it can also be important to provide cues to others as to who is speaking. For agent-agent communication, it is easy to put identifying information in the message channel itself. For humans, however, it may be helpful to provide other cues, such as different voices, and visual cues such as lip movement and gestures for the speaking agent's body or avatar.

1.3 Addressee Recognition

In the two party case, like speaker identification, addressee identification is trivial: whoever is not speaking is the intended recipient of an utterance. In the multi-party case, we must consider *hearers* and *addressees* separately, as discussed above. Hearers of a spoken utterance can be computed by properties such as volume-level of speech, ambient noise, and distance and perceptual abilities of other agents. For agent messages delivered through a router, or other network channels, it may be possible to specify the exact set of receivers of the message.

For calculating the addressee(s) of an utterance several types of information can be used. First, the speaker may directly indicate the addressee using a vocative expression (e.g., calling by name or role). One may also use information included in the content of an utterance, if, e.g., it would be clear that that content would only be addressed to a specific individual. Context is also an important clue – e.g., who had previously spoken or been addressed. If multimodal information is available this can also play an important clue: e.g., gaze or body orientation at a particular individual. Likewise, attention getting or deictic gestures are also clues. If one is the only observable hearer, that can also be a reason to assume the hearer is the addressee. The algorithm used for computing addressees in the MRE project is shown in Figure 1.

1. **If** utterance specifies addressee (e.g., a vocative or utterance of just a name when not expecting a short answer or clarification of type person)
then Addressee = specified addressee
2. **else if** speaker of current utterance is the same as the speaker of the immediately previous utterance
then Addressee = previous addressee
3. **else if** previous speaker is different from current speaker
then Addressee = previous speaker
4. **else if** unique other conversational participant
then Addressee = participant
5. **else** Addressee unknown

Fig. 1. MRE Agent Speech Addressee Identification Algorithm

1.4 Other Participant Roles

In addition to the conversational roles, there are also specific task roles, relating participants to tasks in a variety of ways. In two-party dialogue, typically agents are either performers of a task or those who desire the task to be done, although more complex relationships are possible. For multi-party team situations, such as those in MRE, more complex models are required to support negotiation and team action [4]. We distinguish the agent who will perform a primitive task, from the agent who is *responsible* for a complex task (this agent might perform

all of the sub-actions, or might coordinate a team of actors). Also, some tasks have a *authority*, who can authorize the team-members to carry out the task. This might be different both from the responsible party, the performers of the primitive acts, and agents who actually desire the task to be performed. Agents might also be *guards* for a task, e.g., making sure that it is *not* performed.

Some activities involving dialogue have specific roles, each with designated rights and responsibilities concerning participation in dialogue. This is true even for two party dialogue, such as shopkeeper and buyer, or information seeker and information provider, however much more complex relationships are possible with multiple participants and roles. These can include the ability and length and content of turns, right to assign turns, right to set and change the topic of the conversation. Courtroom dialogue is a striking case with many distinct roles, such as judge, clerk, prosecutor, defense counsel, and witness [8]. Roles may be filled by a single individual, or multiple individuals may fulfil the same role. Likewise, a single individual may play multiple roles.

There are also *social roles* that go beyond a single activity, but structure multiple interactions and tasks. Two types of social roles include status roles (e.g., superior, subordinate, equal, incomparable), and closeness (e.g., friend, comrade, colleague, acquaintance, stranger, opponent, antagonist). These roles will influence the kinds of interaction allowed (e.g., only a superior may give an order to a subordinate), to how likely one will be to adopt the attitudes of another, or comply with their perceived desires. There are also *institutional roles*, such as office in a company, or military rank, defined by the institution.

2 Interaction Management

There are a number of aspects of managing the flow of communication, including the issues of who speaks when, what is the topic under discussion (and how it shifts), and what communicative channels are used (for which topics). Each of these are research topics even for two-party conversation, but become more complex with multiple agents.

2.1 Turn management

There has been a fair amount of work on turn-taking even for two-party dialogue. The basic questions are when to speak and when to stop speaking. Older dialogue systems generally force rigid turn-taking, where one party must wait until the other finishes before speaking. Many more recent systems allow “*barge-in*”, where a human who already understands a system query may provide the answer before the system has finished the utterance. Other systems allow interruptions by both parties, to correct or initiate something new, as well as to respond to the current utterance. Speakers can give verbal and non-verbal signals of continuation or imminent termination of the turn. Speakers use prosody, sentence structure, filled pauses (e.g., “uhhh”), as well as gaze and gesture. Turn-taking

can be modelled using these cues as well as timing information to recognize turn-taking acts [9] such as **take-turn**, **release-turn** and **keep-turn**.

In multi-party dialogue turn-taking is more complex, since more agents are available to potentially take the turn. As well as simply more agents competing for the turn, more actions are possible, e.g., assigning the turn to a particular next speaker vs just releasing it to whoever wants to speak next. Likewise, one may need to request the turn in order to be able to take it, especially if one is not already an active participant.

2.2 Channel management

In uni-modal communication systems, such as simple telephone speech systems, channel management is very similar to turn-management, though differences may arise if the communication channel enforces a single communicator at a time (as with half-duplex circuits, or chat systems which allow only one person to type at a time). In multi-channel systems, however, there is an additional issue of which channel to use for which content, as well as the timing of the contributions. Channels can be using the same modality (e.g., a radio with different frequencies, or a chat system with different chat rooms or different communication commands), or different modalities, e.g., in the MRE system, agents can use verbal communication for face to face or radio communications, and can also use gaze and gesture in the visual mode for face to face communications. One could thus use the speech channel as the main communicative mode, while using the visual mode for *backchannels*, indicating attention and understanding.

For multi-party dialogue, one can simultaneously have multiple “main-channels”, e.g., one per topic, one per conversation, or one per set of participants. Thus, one may have simultaneous communication that is not interruption, because of occurring on different channels between different participants.

2.3 Thread/Conversation management

Turn and channel management concern when and where communication take place. Thread management concerns what is being communicated, specifically which topics are discussed when, and how to organize the progression of topics. Traditional models follow a stack-based topic organization [10], in which one can have hierarchical organization of topics, but not parallel topics under discussion at the same time – when one goes back to a previous topic, one should “pop” the current topic from the stack. Even for two-party conversation, this may be too restrictive [11], especially when multiple channels can be used (e.g., many chat systems, in which two people can type simultaneously without seeing the text until one hits return, and topics often proceed in pairs). With multiple participants, it is also much easier to keep multiple topics open, with different sets of participants.

Another issue is that of multiple conversations. Most current dialogue systems are concerned with only a single conversation with a single user. In contrast, many tasks require different periods of communication separated by periods of

task performance or maintenance in which no communication is required. While some of the information that is conveyed during a prior communication episode is maintained by the participants, often the specific dialogue structure such as the turn and topic structure is not preserved. While it maybe be best to model separate conversations even for extended two-party dialogue, it is essential for multi-party dialogue, where multiple groups of participants communicate with different groups, using different media, about different topics. Having multiple conversation models allows each one to have its own structure, which can be simple and independent of the structure of other conversations that might be going on at the same time. For example, in the MRE Bosnia domain, there is usually a main conversation between Lieutenant, Sergeant and sometimes medic, and subordinate conversations between the Lieutenant and other units over the radio, and between the sergeant and troop members on specific tasks. Each conversation has its own starting, body, and ending phases, as well as participant roles. In some circumstances, especially when multiple participants are part of a conversation, participants can dynamically enter and leave a conversation while it is ongoing. In more complex situations, such as cocktail party conversation, conversations can also split and merge dynamically.

Sometimes multiple conversations are not completely independent. This occurs especially when they share a participant, so that different conversations must compete for attention of the participant. Sometimes topics are linked as well. One conversation might be dependent on another, E.g., if agent *A* asks agent *B* a question in conversation *m*, and then *B* must query agent *C* in conversation *n* in order to reply to *A*. In this case conversation *m* is dependent on conversation *n*, at least for that content. Sometimes conversations are not dependent, but influenced by another. E.g., when participants overhear another conversation and take up the same topic (or comment on the other conversation in some way).

When multiple threads are going on at the same time, it can be tricky to determine which thread a particular utterance belongs to. For the two-party, single conversation case, one can usually rely on topical coherence and cue phrases to determine whether the current utterance continues an existing thread, ends a thread, or begins a new one (and at which level of structure). With multiple participants and multiple conversations which may share participants, the problem becomes more difficult. One can use a number of relationships to try to match the utterance to the proper conversation. There may be a connection between a conversation and a channel, in that case observing the utterance on that channel may help determine the conversation. Likewise, there is a relationship between the addressee and the conversation. As in Figure 1, where knowledge of the conversation was used to help predict the addressee, knowledge of the addressee can point to a conversation containing that addressee as a participant. There is also a relationship between topics and conversations. Identifying the topic of an utterance may help determine which conversation it belongs to, and vice versa.

2.4 Initiative Management

Initiative (or control) [12–16], concerns which agent is currently setting the agenda for topics of discussion. If one agent has the initiative, then another agent does take turns, but only to react to what was said, not to start new topics. Two-party dialogue systems are traditionally either user-initiative (such as question answering systems, where a user may pose a query, and the system consults a database and provides an answer) or system-initiative, in which the system asks a series of queries to specify the parameters for a service request. More recently, *mixed-initiative* systems allow user and system to both take the initiative at different points. E.g., system can take the initiative when there are problems in communication, to direct toward possible solutions, and human can take control to more efficiently provide known information.

In multi-party dialogues, initiative is less symmetric than two party dialogues for equivalent tasks [17]. Thus, the more participants in a conversation, the less likely it will be that each participant has an equal amount of initiative. Team leaders tend to develop, either formally, or informally, who structure the interaction. Other kinds of initiative are also possible, e.g., cross-initiative, where a responder does not take initiative herself, but redirects it to a third party (who might not even have been active), or in which a third party interjects. There are also issues of cross-conversation initiative, e.g. in the case of one conversation being dependent on another, the initiative-holder of one conversation is really taking direction from someone else in another conversation.

2.5 Attention management

Attention is mostly assumed to be always present for most single-user, single-system dialogue systems. Even when attention is explicitly modelled, it is usually a binary decision of either being on the conversation and other participant, or elsewhere. In multi-party, multi-conversation situations, however, a much more detailed model of attention is required. An attention model can be used to summon others into a new or existing conversation, and can model which conversation each participant is attending to.

3 Grounding and Obligations

Much of the local content of dialogue can be modelled using notions like obligations and grounding [9, 18–24]. These models become more complex when considering the multiparty case.

Grounding is the process of adding to the common ground between participants in conversation [24]. The grounding model in [9, 19, 25] consisted of a structure of *Common ground units*, (CGUs) each of which contains material that is added to the common ground together. Each CGU has a unique initiator, responder, contents and state. The state is calculated using a finite state automaton, updated by *grounding acts* performed on the CGU. States include

those in which the contents are grounded and ungroundable, as well as intermediate states in which an acknowledgement or repair is needed from one party or another. By recognizing grounding units and the CGUs that they construct and add to, a computational agent is able to model and participate in the grounding process.

In the MRE project, this model has been used in multiparty conversation, but only in cases in which there is a single initiator and responder of a particular CGU. For the more general case, in which there are multiple addressees, it is less clear what the proper grounding model *should* be. One option is to allow any of the addressees to acknowledge for the contents to be considered grounded. The problem is that this may lead to overly optimistic [26] estimations of common ground, where some agents did not in fact understand or possibly receive the communications. The pessimistic extreme is to require evidence of understanding from each addressee. While this is safer, it seems somewhat unrealistic when many of the addressees are human. Some sort of middle-ground is also possible, requiring an amount of evidence that is more than a single acknowledgement from one agent, but less than a separate acknowledgement from each agent.

Another interesting issue is grounding across conversations. E.g., if *A* asks *B* a question and observes *B* asking the same question to *C* (whether in the same conversation or a different one), *A* has evidence that *B* has understood the question, even though *B* has not yet responded to *A*.

Multiple addressees also present a challenge for models of obligation. The model of discourse obligations presented in [19–21] takes one of the main effects of utterances like requests and questions to be an obligation to perform some action such as addressing the request (by performing the requested action, accepting or rejecting the request, or other negotiating or explaining move). When there are multiple addressees, however, it is not so clear what the status of these obligations are. Does every addressee have a personal obligation? Is there an indefinite obligation assigned to the group, that can be satisfied by any member performing an obligation-relieving action? In the case of this indefinite obligation, what is it that motivates any particular agent to act?

Also there is the issue of transfer of obligation. To take the example given above, where *B* redirects *A*'s question to *C*, if this is done in the presence of *A*, does *B* still have the obligation? Whether or not *B* still holds the obligation, does *C*'s response in *A*'s presence relieve *B* of this obligation? Can another party, say *D* relieve the obligation by providing an answer even when not addressed? The answers to some of these questions depend on the particular type of activity. For instance, if the purpose of *A*'s question is to solicit information, and *C* or *D* are trustworthy, probably no more action is required of *B*. On the other hand, if it is a classroom situation, where *A* is asking the question not so much to find out the answer, but to determine whether *B* knows it, then *B*'s redirect to *C* and *D*'s spontaneous reply would be out of place, and perhaps subject to sanctions.

In some cases, multi-party dialogue can actually make the theoretical models of dialogue clearer rather than obscuring them. A case in point is an account of what motivates agents to answer questions. As described above, one model that

has been used in some dialogue systems takes obligations as the motivation; the systems are designed to track obligations and then use these to motivate performing answers. An alternate model has been to use dialogue structural considerations, such as *Questions Under Discussion* (QUD), based on work by Ginzburg [27] to model question answering. When a question is asked, it gets added to the QUD, which in turn licenses answers to the question (including elliptical short answers). Both approaches were used in the TRINDI project [28, 29]. The GoDiS system [30] uses a QUD structure, while the EDIS system [31], uses the obligation approach. For simple two-party information-seeking domains such as Autoroute [29], there is little to choose between these two accounts. Both do an adequate job of representing questions, answers, intermediate states, and observation of lack of answers or other responses.

However we can see that there are really some distinct functions, as pointed out in [32]. QUD represents information about what would count as an answer, while obligations represent who should/must answer. Both reflect on the question of when the answer should occur. Obligations may specify time-limits on the answer. QUD, on the other hand will allow one to track when a particular utterance could be understood as an answer to that question. E.g., if an intervening question of a similar type is asked after the original question, a new utterance may be taken as an answer to the second rather than first question. In the MRE dialogue model, we represent both QUD and obligations. The former is part of the conversation structure of a specific conversation, while the latter (if grounded) is a property of the social state between agents. Thus an obligation might be introduced by a question in one conversation, and relieved in another conversation. The form of the answer depends on the QUD structure, however. If a question is not on QUD in the current conversation, then the question must be reintroduced before answering, or at least the answer must be given with sufficient clarity to accommodate the question [33].

4 Conclusions

In this article we have examined a number of issues in dialogue management for how they scale when moving from a two-participant model to a multi-participant model. Two obvious choices are available for multi-party models. One is to treat multiparty conversation as a set of pairs of two-party conversations. While this has the advantage of simplicity and using existing models, it is less than satisfactory in some cases. In the worst case, one will still need to move beyond the two party case in order to arbitrate between the multiple interactions, e.g. *A* with *B* and *A* with *C*. In some cases this will be more complex than changing the model to allow multiple participants. In some cases, we can see two-party dialogue as a special simple case of multiparty dialogue.

Dialogue system evaluation is also a difficult subject even for two-party dialogue. There are no universally agreed on metrics, due in large part to the very different types of tasks that dialogue systems are used for. Still, there are some general themes for evaluation, including task success, naturalness of interaction,

user satisfaction, and efficiency. Some of these can be applied to the multi-party case, but the metrics become more difficult to calculate. E.g., for efficiency does one count real-time, or total agent time? One might count only a human's time, but what if there are multiple humans? Similar issues exist for other issues - how does one count naturalness when some agents communicate fairly naturally but others don't?

We are as yet only in the beginning stages of modelling multi-party dialogue, with few applications and very few implemented systems. The requirements will surely increase, however, as more societies of agents and people interact in more fluid ways.

Acknowledgements

This article was first presented as an invited talk to the 4th Sigdial workshop on discourse and dialogue, July 2003 in Chiba, Japan. The work described in this paper was supported by the Department of the Army under contract number DAAD 19-99-D-0046. Any opinions, findings and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the Department of the Army. The author would like to thank many colleagues for helping develop the ideas and intuitions presented in this paper.

References

1. Traum, D.R., Rickel, J.: Embodied agents for multi-party dialogue in immersive virtual worlds. In: Proceedings of the first International Joint conference on Autonomous Agents and Multiagent systems. (2002) 766–773
2. Rickel, J., Marsella, S., Gratch, J., Hill, R., Traum, D., Swartout, W.: Toward a new generation of virtual humans for interactive experiences. *IEEE Intelligent Systems* **17** (2002)
3. Swartout, W., Hill, R., Gratch, J., Johnson, W., Kyriakakis, C., Labore, K., Lindheim, R., Marsella, S., Miraglia, D., Moore, B., Morie, J., Rickel, J., Thiebaux, M., Tuch, L., Whitney, R., Douglas, J.: Toward the holodeck: Integrating graphics, sound, character and story. In: Proceedings of 5th International Conference on Autonomous Agents. (2001)
4. Traum, D., Rickel, J., Marsella, S., Gratch, J.: Negotiation over tasks in hybrid human-agent teams for simulation-based training. In: In proceedings of AAMAS 2003: Second International Joint Conference on Autonomous Agents and Multi-Agent Systems. (2003) 441–448
5. Marsella, S., Gratch, J.: Modeling coping behavior in virtual humans: Don't worry, be happy. In: In proceedings of AAMAS 2003: Second International Joint Conference on Autonomous Agents and Multi-Agent Systems. (2003)
6. Clark, H.H.: Using Language. Cambridge University Press, Cambridge, England (1996)
7. Brunelli, R., Falavigna, D.: Person identification using multiple cues. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **17** (1995) 955–966

8. Martinovski, B.: The Role of Repetitions and Reformulations in Court Proceedings – a Comparison of Sweden and Bulgaria. PhD thesis, Göteborg University: Department of Linguistics (2001)
9. Traum, D.R., Hinkelmann, E.A.: Conversation acts in task-oriented spoken dialogue. *Computational Intelligence* **8** (1992) 575–599 Special Issue on Non-literal language.
10. Grosz, B.J., Sidner, C.L.: Attention, intention, and the structure of discourse. *Computational Linguistics* **12** (1986) 175–204
11. Walker, M.A.: Limited attention and discourse structure. *Computational Linguistics* **22** (1996) 255–264
12. Whittaker, S., Stanton, P.: Cues and control in expert-client dialogues. In: *Proceedings ACL-88.* (1988) 123–130
13. Walker, M.A., Whittaker, S.: Mixed initiative in dialogue: An investigation into discourse segmentation. In: *Proceedings ACL-90.* (1990) 70–78
14. Novick, D.: Control of Mixed-Initiative Discourse Through Meta-Locutionary Acts: A Computational Model. PhD thesis, University of Oregon (1988) also available as U. Oregon Computer and Information Science Tech Report CIS-TR-88-18.
15. Chu-Carroll, J., Brown, M.K.: Tracking initiative in collaborative dialogue interactions. In: *Proceedings of the Thirty-Fifth Meeting of the Association for Computational Linguistics, Association for Computational Linguistics* (1997) 262–270
16. Ishizaki, M.: Mixed-Initiative natural Language Dialogue with Variable Communicative Modes. PhD thesis, University of Edinburgh (1997)
17. Ishizaki, M., Kato, T.: Exploring the characteristics of multi-party dialogues. In: *COLING-ACL.* (1998) 583–589
18. Allwood, J.: Obligations and options in dialogue. *Think Quarterly* **3** (1994) 9–18
19. Traum, D.R.: A Computational Theory of Grounding in Natural Language Conversation. PhD thesis, Department of Computer Science, University of Rochester (1994) Also available as TR 545, Department of Computer Science, University of Rochester.
20. Traum, D.R., Allen, J.F.: Discourse obligations in dialogue processing. In: *Proceedings of the 32nd Annual Meeting of the Association for Computational Linguistics.* (1994) 1–8
21. Poesio, M., Traum, D.R.: Towards an axiomatization of dialogue acts. In: *Proceedings of Twendial'98, 13th Twente Workshop on Language Technology: Formal Semantics and Pragmatics of Dialogue.* (1998)
22. Dignum, F., van Linder, B.: Modeling social agents: Communication as action. In Müller, J.P., Wooldridge, M.J., Jennings, N.R., eds.: *Intelligent Agents III — Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96).* Lecture Notes in Artificial Intelligence. Springer-Verlag, Heidelberg (1996)
23. Cahn, J.E., Brennan, S.E.: A psychological model of grounding and repair in dialog. In: *Working Notes AAAI Fall Symposium on Psycholiogical Models of Communication in Collaborative Systems.* (1999)
24. Clark, H.H., Schaefer, E.F.: Contributing to discourse. *Cognitive Science* **13** (1989) 259–294
25. Poesio, M., Traum, D.R.: Conversational actions and discourse situations. *Computational Intelligence* **13** (1997)
26. Poesio, M., Cooper, R., Larsson, S., Matheson, C., Traum, D.: Annotating conversations for information state update. In: *Proceedings of Amstelogue'99 workshop on the semantics and pragmatics of dialogue.* (1999)
27. Ginzburg, J.: Interrogatives: Questions, facts and dialogue. In Lappin, S., ed.: *The Handbook of Contemporary Semantic Theory.* Blackwell, Oxford (1996)

28. Larsson, S., Traum, D.: Information state and dialogue management in the TRINDI dialogue move engine toolkit. *Natural Language Engineering* **6** (2000) 323–340 Special Issue on Spoken Language Dialogue System Engineering.
29. Traum, D., Bos, J., Cooper, R., Larsson, S., Lewin, I., Matheson, C., Poesio, M.: A model of dialogue moves and information state revision. Technical Report Deliverable D2.1, Trindi (1999)
30. Bohlin, P., Cooper, R., Engdahl, E., Larsson, S.: Information states and dialogue move engines. In: Proceedings of the IJCAI99 workshop: Knowledge And Reasoning in Practical Dialogue Systems. (1999) 25–31
31. Matheson, C., Poesio, M., Traum, D.: Modelling grounding and discourse obligations using update rules. In: Proceedings of the First Conference of the North American Chapter of the Association for Computational Linguistics. (2000)
32. Traum, D.: Semantics and pragmatics of questions and answers for dialogue agents. In: proceedings of the International Workshop on Computational Semantics. (2003) 380–394
33. Larsson, S.: Issue-based Dialogue Management. PhD thesis, Göteborg University (2002)